

REMARKS

Claims 1-20 were presented for examination and were pending in this application. In the latest Office Action, claims 1-8 were rejected, and claims 9-20 were withdrawn from consideration as directed to a nonelected invention. With this amendment, claims 1-3 and 5-8 are amended, and the withdrawn claims 9-20 are canceled. On the basis of the following remarks, consideration of this application and allowance of all pending claims are requested.

Several objections to the claims were raised in the Office Action. The claims have been amended to address the examiner's objections, as suggested in the Office Action.

Claims 1-8 were rejected as anticipated by U.S. Patent No. 6,246,692 to Dai et al. Applicant respectfully asserts that the claims, as amended, are novel and patentable over Dai.

The amended claims recite a network system for interconnecting a set of packet-switching network elements. The claimed network system comprises a set of nodes that interface with the packet-switching network elements, where the nodes are connected to each other by variable capacity connections. In each connection, data are transported from a source node to a destination node, and the connection has a capacity and a traffic load. In the claimed network system, the capacity of each connection is controlled based at least in part on the traffic loads associated with the connections configured to transport data to the same destination node. Dai does not anticipate the claimed network system for a number of reasons, at least some of which are provided below.

"network system comprising a set of nodes"

The claims, as amended, recite a network system that interconnects packet-switching network elements and optimizes communications among a plurality of nodes. In this

amendment, the previously-recited “interface units” have been changed to “nodes.” This amendment finds support in the specification (e.g., at p. 11, 13-14, and FIGS. 2B, 2C, and 4), which explains that the interface units, in a preferred embodiment, are distinct nodes and may be separate physical nodes or physically integrated (i.e., clustered) with each other or with external network elements. As claimed, a node is an independent network element. The claimed network system comprising multiple nodes is thus contrasted with Dai, which describes a system for backplane bandwidth allocation among interface cards of a single node.

The system described and illustrated in Dai is internal to a single node, i.e., element 10 in FIG. 1, the “packet switching fabric.” (See, e.g., Dai, col. 1, lines 65-67; col. 2, lines 30-34.) Because Dai’s switching devices are controlled by a local management device of their common single node, they are not separate nodes as in the claimed network system. Dai’s switching devices thus do not provide communications among nodes, since Dai only discloses a single node. It is further noted that the background section of Dai discusses the switching fabric as a system with a backplane and network interface cards (col. 1, lines 50-52) and, in the context of discussing its objectives, mentions the source-destination data links as “internal” to the fabric (col. 1, lines 65-67; col. 2, lines 30-34). Therefore, the scope of Dai for bandwidth allocation is limited to applications internal to a single node, such as an Ethernet switch, controlled by its single management device and CPU.

This distinction is also important in at least one respect because each of the claimed nodes can be located at arbitrary distances from one another, whereas Dai’s switching devices are located within a single node. This is because Dai’s switching fabric requires a single dedicated management device and a CPU that are required to control the switch devices over parallel control buses. It would not be physically feasible to control distant nodes via Dai’s

control ring due to the unavoidable differences in signal propagation delays among the individual wires that form the parallel signal path of Dai's control ring. (In practice, the maximum workable length of such parallel signals in high-speed nodes is limited to a few centimeters.)

"the capacity of each connection controlled . . . based at least in part on the traffic loads associated with the connections configured to transport data to that destination node"

The claimed system enables a destination node to control how to optimize its connection capacity allocation among a set of concurrent connections from multiple source nodes at the same time and based on the traffic loads toward that destination node from multiple source nodes. In contrast, Dai only allows the destination device to select from which single source device of its local node to receive data at a given time. Dai does not allow for multiple concurrent connections to a single destination node, wherein "each one of the connections [are] configured to transport data from its source node to its destination node." In Dai, the bandwidth of a source-destination channel when the channel is transferring data is commensurate to the capacity of the destination port associated with that channel. Therefore, Dai does not enable a destination node of a network system to control simultaneously the capacities of multiple connections based on traffic loads associated with the connections from multiple source nodes to that destination node. Accordingly, Dai does not disclose the claimed feature of controlling the capacity of each connection from its destination node "based at least in part on the traffic loads associated with the connections configured to transport data to that destination node."

In the Office Action, the examiner suggested that Dai discloses at column 6, lines 11-13, "that the amount of bandwidth (capacity) allocated for each source-destination channel is commensurate with the network link capacity of a destination network port (destination interface

unit).” But as described above, Dai does not disclose a system in which capacity of each connection is controlled from the destination node. The passage cited by the examiner merely states that the allocated bandwidth for a given channel, when such a channel exists, is equal to the capacity of the destination network port. It does not disclose that the “capacity of each connection [is] controlled from its destination node,” as the switching fabric of Dai does not support more than one channel per a destination network port at any given time (see, e.g., Dai, col. 4, lines 11-12); thus, at any given time, there is a single destination device corresponding to a given source device of the switching fabric. In Dai, therefore, a destination device cannot have more than one channel to it from any of the source devices of the switching fabric at any given time.

“providing a connection of variable capacity to the other nodes of the network system”

Moreover, Dai is not a network system in which each node “provid[es] a connection of variable capacity to the other nodes of the network system,” as claimed. As explained in Dai, col. 6, lines 11-13, the bandwidth of a source-destination device channel, when set up for data transfer, is commensurate with the network link capacity of the destination network port. In other words, a single channel within the switching fabric of Dai fully consumes the capacity of its associated destination network port. Accordingly, the devices of Dai cannot concurrently have connections to multiple other devices of the system because a destination device of Dai can receive a channel from only one source device at any given time.

For at least the foregoing reasons, Dai does not anticipate claim 1 or any of claims 2-8, which depend from claim 1. Moreover, claims 2-8 are further patentable for their additionally recited dependent limitations.

For instance, claim 2 further recites that “the system is configured to set the capacity of a connection to zero when the connection has no traffic load associated therewith and traffic loads associated with other connections to the same destination node cumulatively exceed a predefined limit.” This feature is not disclosed in Dai, nor would it be possible in Dai’s system. In Dai’s system, a channel across the fabric backplane must be distinctly set up before data can be transferred, and the backplane bandwidth reserved for a channel must likewise be released after each data burst using specific messaging procedures. (Dai, col. 3, lines 35-38; col. 6, lines 5-11; and FIGS. 8B and 8E.) Furthermore, these channel set-up, bandwidth reservation, and bandwidth release procedures must be performed individually for each packet to be connected through the node. (Dai, col. 3, lines 50-52.) In Dai, therefore, a source-destination channel does not exist when its bandwidth is zero, since channels per Dai need to be specifically set up to allocate bandwidth for such channels. Claim 2 is further patentable over Dai for at least this reason.

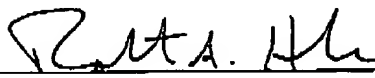
Because Dai’s system is a single node, Dai does not anticipate the claimed features further recited in claim 5 (a multi-node network system wherein some of the nodes are physically located at a single physical platform or attached to a single chassis) or claim 6 (a network system with some of its nodes integrated into their associated packet-switching network elements). In addition, it would not be physically feasible to control distant nodes using Dai’s parallel signal paths made of collections of individual bit-specific wires (i.e., Dai’s data and control rings 16, 18). Therefore, Dai cannot anticipate the claimed features further recited in claim 7 (wherein the

system is at least in part a sub-network of a multi-use or public network). Finally, because Dai does not disclose a multi-node network system for interconnecting a set of packet-switching network elements, as described above, it also fails to anticipate claim 8 (where a packet-switching network element is network system according to claim 1).

Based on the foregoing, the application is in condition for allowance of all claims, and a Notice of Allowance is respectfully requested. If the examiner believes for any reason direct contact would help advance the prosecution of this case to allowance, the examiner is encouraged to telephone the undersigned at the number given below.

Respectfully submitted,
MARK SANDSTROM

Dated: January 18, 2006

By: 
Robert A. Hulse, Reg. No. 48,473
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2444
Fax: (415) 281-1350